

Diffusive behavior and the modeling of characteristic times in limit order executions

Zoltán Eisler,^{1,2,*} János Kertész,^{2,3} Fabrizio Lillo,^{4,5,6} and Rosario N. Mantegna^{4,6}

¹*Science & Finance, Capital Fund Management, Paris, France*

²*Department of Theoretical Physics, Budapest University of Technology and Economics, Budapest, Hungary*

³*Laboratory of Computational Engineering, Helsinki University of Technology, Espoo, Finland*

⁴*Dipartimento di Fisica e Tecnologie Relative, Università di Palermo, Viale delle Scienze, I-90128, Palermo, Italy*

⁵*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA*

⁶*CNR-INFM, Unità Operativa di Roma, Centro di Ricerca e Sviluppo SOFT, Roma, Italy*

(Dated: December 21, 2008)

We present an empirical study of the first passage time (FPT) of order book prices needed to observe a prescribed price change Δ , the time to fill (TTF) for executed limit orders and the time to cancel (TTC) for canceled ones in a double auction market. We find that the distribution of all three quantities decays asymptotically as a power law, but that of FPT has significantly fatter tails than that of TTF. Thus a simple first passage time model cannot account for the observed TTF of limit orders. We propose that the origin of this difference is the presence of cancellations. We outline a simple model, which assumes that prices are characterized by the empirically observed distribution of the first passage time and orders are canceled randomly with lifetimes that are asymptotically power law distributed with an exponent λ_{LT} . In spite of the simplifying assumptions of the model, the inclusion of cancellations is enough to account for the above observations and enables one to estimate characteristics of the cancellation strategies from empirical data.

Keywords: Econophysics, Limit order book, First passage time, Brownian motion, Time to fill

I. INTRODUCTION

Understanding the market microstructure is crucial for both theoretical and practical purposes [1]. On double auction markets the limit order book contains most of the information about the market microstructure and price discovery. Recently there has been considerable effort to investigate limit order book dynamics. Empirical studies [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17] have been devoted to the search for the key determinants of price formation, the trading process and market organization. A large number of papers have focused on modeling the limit order book with [18, 19, 20, 21, 22] or without [23, 24, 25, 26] dynamics. Market microstructure studies consider a large number of aspects of the price discovery mechanism and these studies can greatly contribute to the success of the modeling of financial markets. The market mechanism, along with the complex interactions among market participants results in the emergence of a collective action of continuous price formation. Some of the studies have used an agent based modeling approach. Examples are market models described in terms of agents interacting through an order book based on simple rules [27, 28] and models where the assumptions about the trading strategies are kept as minimal as possible [14, 19]. One of the most striking findings was that even if trends and investor strategies are neglected, purely random trading may be adequate to describe certain basic properties of the order book [13].

Most of the above papers focus on limit order executions, and very few deal with cancellations, even though

the frequency of the two outcomes is comparable [9]. The uncertainty of execution represents a primary source of risk [29]. Another major risk factor is adverse selection, also known as "pick-off" risk. This risk is associated with the waiting time until order execution. During this period those with excess information can take advantage of the liquidity provided by the limit orders of less informed traders, and hence it is important to accurately quantify these waiting times. Lo et al. [9] apply survival analysis to limit order data, and they find that the time between order placement and execution is very sensitive to the limit price, but not to the volume of the order. They also investigate the dependence on further explanatory variables such as the bid-ask spread and the volatility. The dynamics of the limit order book has also been investigated by using a joint model of executions and cancellations in a framework of competing risks¹. Within this approach Hollifield et al. [16], by using observations on order submissions and execution and cancellation histories, estimate both the distribution of traders' unobserved valuations for the stock and latent trader arrival rates. Chakrabarty et al. [29] show that executions are more sensitive to price variation and less to volume variation than cancellations. This last work also analyzes the

¹ The notion of competing risks applies to problems where one deals with several "risks", i.e., random events, of which only the first one can be observed [30]. For example, limit orders are either executed or canceled and both events can be modeled by some random process. If an order gets canceled, one can no longer directly observe what time it would have been eventually executed, and vice versa. Thus it is not possible to independently estimate either process without a bias, if one simply ignores information from the other one.

*Electronic address: eisler@maxwell.phy.bme.hu

relationship between execution time and market depth.

In this paper we aim to go a step further, and combine the framework of competing risks with random walk theory. In particular, we analyze the difference observed between the time to fill a limit order, which is the time one had to wait before a limit order was executed, and the first passage time [31], i.e., the time elapsed between an initial instant and the time when the transaction price crosses a given predefined threshold. In addition, the largest difference between our approach and most previous studies (e.g., Refs. [9, 29]) is that while those placed more emphasis on the typical values of execution and cancellation times, we will concentrate on the accurate description of the rare events, and the related asymptotic tail behavior of the distributions.

We observe that for a fixed price change the first passage time distributions of transaction price, best bid and best ask are quite well described asymptotically by the theoretical form expected for a Markov process with symmetric jump length distribution (including Brownian motion) [31, 32]. The empirical time to fill of executed orders is smaller than the first passage time. We attribute this difference to canceled and expired orders. We propose a simple competing risks model, where limit orders are removed from the order book when either of two events happens: (i) when they are executed, this is modeled as the first time when the transaction price reaches the limit price, (ii) or when they are canceled, the time horizon of cancellations is modeled as a random process that is independent from price changes. In this framework we are able to predict constraints about the tail behavior of the time to fill and time to cancel probability densities. Our model also allows us to estimate the distribution of the time horizons of the placed limit orders. We show that the assumption of independence between the price changes and order cancellations, while it is a large simplification compared to real data, does not affect our conclusions significantly.

The paper is organized as follows. In Section II we describe the investigated market and the variables of interest. In Section III we study the first passage time and in Section IV the time to fill and the time to cancel. Section V describes a simple limit order model and Section VI is devoted to testing the model empirically. Section VII extends the result to limit orders placed inside the spread. Section VIII discusses the validity of the assumptions and summarizes the results. Finally, in the Appendices we show that the results are unchanged if time is measured in transactions. Then we present a critical discussion of the fitting procedure we used to estimate the tail behavior of the time to fill and time to cancel distributions.

II. THE DATASET

The empirical analysis presented in this study is based on the trading data of the electronic market (SETS) of

London Stock Exchange (LSE) during the year 2002. These data can be purchased directly from the London Stock Exchange. We investigate 5 highly liquid stocks, AstraZeneca (AZN), GlaxoSmithKline (GSK), Lloyds TSB Group (LLOY), Shell (SHEL), and Vodafone (VOD). Opening times of LSE are divided into three periods. The intervals 7:50–8:00 and 16:30–16:35 are called the opening and the closing auction, respectively. These follow different rules and thus also observe different statistical properties than the rest of the trading. Therefore we discarded limit orders placed during these times, and focused only on the periods of continuous double auction during 8:00–16:30. We also removed limit orders that were placed during 8:00–16:30 but were canceled (or expired) during the opening/closing auctions. We measure time intervals in trading time, i.e., we discard the time between the closing and the opening of the next day.² Finally, whenever we refer to prices we exclude all transactions that were executed on the SEAQ market³ and not in the limit order book.

We denote the best bid price⁴ by $b(t)$, the best ask price by $a(t)$ and the bid-ask spread is $s(t) = a(t) - b(t)$. Except for very special cases, there are already other limit orders waiting inside the book when one wants to place a new one. Let $b(t) - \Delta$ denote the price of a new buy limit order, and $a(t) + \Delta$ the price of a new sell limit order. Orders placed exactly at the existing best price correspond to $\Delta = 0$, orders placed inside the spread have $\Delta < 0$, while $\Delta > 0$ means orders placed "inside the book". It is possible to have so called crossing orders with such large negative values of Δ that they cross the spread, i.e., $\Delta < b(t) - a(t)$. These orders can be partially or fully executed immediately by limit orders from the other side of the book. Since a trader would place a crossing limit order to execute (at least part of) it immediately, we will not consider them as limit orders in our analysis.

Any limit order which was not executed can be canceled at any time by the trader who placed it. The order can also have a predetermined validity after which it is automatically removed from the book, this is called expiry. We will not distinguish between these mechanisms and we will call both of them cancellation. Throughout the paper we will use ticks as units of price and all loga-

² In our analyses, we removed the data of trading on September 20, 2002. This is because on that day very unusual trading patterns were observed, including an anomalous behavior of the bid-ask spread.

³ Many studies refer to this colloquially as the "upstairs" market.

⁴ In most of the literature the logarithm of the price is modeled, while throughout the paper we intentionally use price itself. Our study is concerned with very small price changes on the order of the spread, when there is little difference between the two approaches. In our case it is important to keep bare prices, as stocks have a finite tick size (minimal price change). Taking bare prices enables us to classify the orders into discrete categories by price difference. The size of ticks depends on the stock, the possible values are 1/4, 1/2 or 1 penny.

rithms are 10-base.

III. THE FIRST PASSAGE TIME

Let the latest transaction price of an asset at time $t_0 = 0$ be S_0 . The first passage time [31] of price through a prescribed level $S_0 + \Delta$ with some fixed $\Delta > 0$ is defined as the time t of the first transaction when $S(t) \geq S_0 + \Delta$. Similarly we can determine the first time after $t_0 = 0$ when the transaction price was below or equal to $S_0 - \Delta$ and we will consider this time as another, independent observation of t . We will call the distribution of the quantity t the first passage time distribution to a distance Δ , and denote it by $P_{\text{FPT};\Delta}(t)$.

Such first passage processes have been studied extensively [33]. For simplicity we will restrict ourselves to driftless processes. This is justified, because in real data for time horizons t of up to a day the drift of the prices is negligible. This means that the ratio $|\mu|\sqrt{t}/\sigma$ is small (it is always less than 10^{-1} in our dataset), where μ is the mean price change over unit time, and σ is the standard deviation of price changes during a unit time (i.e., the volatility). Throughout the paper we use real time⁵.

For the following analysis of empirical data, it is useful to review the first passage time distribution for Brownian motion without drift. This is can be written as [31]

$$P_{\text{FPT};\Delta}(t) = \frac{\Delta}{\sqrt{2\pi\sigma^2}} t^{-3/2} \exp\left(-\frac{\Delta^2}{2\sigma^2 t}\right), \quad (1)$$

which is the fully asymmetric 1/2-stable distribution. For any fixed Δ the asymptotics for long times is

$$P_{\text{FPT};\Delta}(t) \propto t^{-3/2}. \quad (2)$$

A recent study [32] has clarified that this asymptotic behavior is valid not only for Brownian motion but also for any Markov process with symmetric jump length distribution.⁶ Of course, real price changes are not described by continuous values, and transactions and order submissions are also separated by finite waiting times,

which a continuous time random walk formalism could take into account [35, 36]. However, in this paper we are interested in time intervals much longer than these waiting times, so the discrete aspects of the dynamics are negligible. Thus, we will model prices as if they varied continuously in time.

Let us now investigate empirically the first passage time behavior. The first passage time distribution for the transaction price, bid and ask when $\Delta = 1$ tick is shown in Fig. 1 for the stock GSK. The distribution is obtained by sampling the first passage time at each second. One can see that there are no significant differences in the behavior of the three prices. Qualitatively, the distribution is similar to Eq. (1), and the long time asymptotic of real data seems to decay approximately as $t^{-3/2}$. For times shorter than 1 minute the curves significantly deviate both from the power law behavior and from the prediction of Eq. (1). We choose to fit the first passage time distribution with the function

$$P_{\text{FPT};\Delta}(t) = \frac{C t^{-\lambda_{\text{FPT}}}}{1 + [t/T_{\text{FPT}}(\Delta)]^{-\lambda_{\text{FPT}} + \lambda'_{\text{FPT}}}}. \quad (3)$$

This form, that we will use to fit also the other distributions introduced below, is characterized by two power law regimes. Normalization conditions of Eq. (3) imply that $\lambda_{\text{FPT}} > 1$ and $\lambda'_{\text{FPT}} < 1$. For $t \ll T_{\text{FPT}}(\Delta)$ it is $P_{\text{FPT};\Delta}(t) \propto t^{-\lambda_{\text{FPT}}}$, whereas for $t \gg T_{\text{FPT}}(\Delta)$ it is $P_{\text{FPT};\Delta}(t) \propto t^{-\lambda_{\text{FPT}}}$. We will discuss the motivations for choosing this form in Section IV and in the Appendix.

Table I contains the fitted parameters λ_{FPT} , λ'_{FPT} , and $T_{\text{FPT}}(\Delta)$ for $\Delta = 1, \dots, 4$ ticks. The difference between the actual values of λ_{FPT} and $3/2$ from Eq. (2) is small. Systematic deviations due to clustered volatility or the fluctuations of trading activity could not be identified. For example, the asymptotic shape of the distribution does not change, even if time is measured in transactions instead of seconds (see Appendix A).

The observation that $\lambda_{\text{FPT}} < 2$ implies that the theoretical mean and standard deviation of the first passage time distribution are infinite. Thus one should be careful with the interpretation of means calculated from finite samples. Throughout the paper we will rely on the determination of quantiles (e.g., the median) instead, which are always well-defined regardless of the shape of the distribution.

The inset of Fig. 1 shows the median first passage time as a function of Δ for the five investigated stocks. The behavior is not exactly quadratic (Δ^2) as one would expect from Eq. (1). If prices followed a Brownian motion, the q -th quantile (T_q) of the first passage time distribution would be

$$T_q = \frac{\Delta^2}{2\sigma^2[\text{erfc}^{-1}(q)]}, \quad (4)$$

where the median ($M[\text{FPT}]$) corresponds to $q = 0.5$. In reality, the power law behavior with Δ is less evident, as shown by the inset of Fig. 1. Assuming a behavior

⁵ We repeated the statistical analysis with transaction time and observed a similar power law decay of the first passage time for large times. The value of the power law exponent turns out to be different for real time analysis and transaction time analysis. See Appendix A for details.

⁶ This result is consistent with the Sparre-Andersen theorem [33]. Alternative descriptions obtained for the asymptotic time dependence of the FPT of Lévy flights which were hypothesizing a dependence of the distribution exponent from the index of the Lévy distribution have missed the fact that the method of images, which is extremely powerful in Gaussian diffusion, fails for Lévy flight processes [32]. The behavior is of course more complex in the case of Lévy random processes described by using a subordination scheme. In these cases the asymptotic behavior of first passage time depends on the complete properties of the subordination procedure [34].

$P_{\text{TTF}}(t)$, $P_{\text{TTFF}}(t)$, and $P_{\text{TTC}}(t)$, respectively.

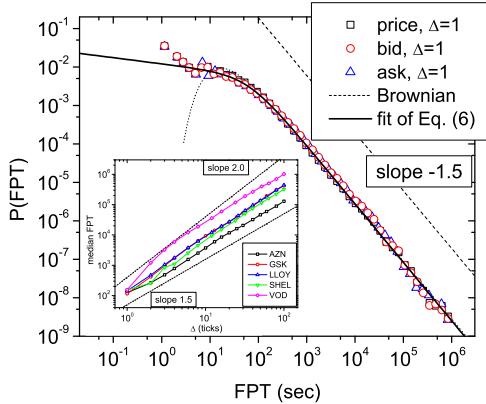


Figure 1: First passage time distributions for the price, bid and ask quotes of GlaxoSmithKline (GSK), distance $\Delta = 1$ tick. The dotted line is the first passage time distribution for Brownian motion with volatility $\sigma = 1/7$ penny \times sec $^{-1/2}$. The thick solid line is a fit with Eq. (3) as given in Table I. The inset shows the median first passage time as a function of Δ .

$M[\text{FPT}] \propto \Delta^\eta$ would require an exponent varying between 1.5 and 1.8 depending on the specific stock and the precise range of Δ used for the estimation of η . A similar deviation from the prediction of Brownian motion was reported in Ref. [37] in the analysis of closure index values sampled at a daily time horizon.

There are many differences between real prices and Brownian motion, and the above non-quadratic behavior can come from any of them: the non-Gaussian distribution of returns, the superdiffusivity of price, perhaps both or none. We have performed a series of shuffling experiments and preliminary results support the conclusion that the main role is played by the deviation from Gaussianity. This non-Gaussianity is well documented in the literature down to the scale of single transactions [12]. A similar effect was seen for Levy flights, whose increments are also very broadly distributed, and their value of η can be different from 2, and it is related to the index of the corresponding Levy distribution [38].

IV. TIME TO FILL, TIME TO CANCEL

For an executed order the time elapsed between its placement and its complete execution is called *time to fill*. Orders are often not executed in a single transaction, thus one can also define *time to first fill*, which is the time from order placement to the first transaction this order participates in. Finally, for canceled orders one can define the *time to cancel* which is the time between order placement and cancellation. The distribution of these three quantities will be in the following denoted by

A. Properties of the distributions

As a first characteristic of the order book, we investigate the distribution of time to fill and time to cancel for the stocks in our dataset. Fig. 2 shows these distributions for GlaxoSmithKline (GSK) for different values of Δ . Similarly to the first passage time, we fitted the empirical density with the function

$$P_{\text{TTF};\Delta}(t) = \frac{C't^{-\lambda_{\text{TTF}}}}{1 + [t/T_{\text{TTF}}(\Delta)]^{-\lambda_{\text{TTF}} + \lambda'_{\text{TTF}}}}. \quad (5)$$

This form (5), which we used to fit the FPT in the previous Section, is different from the more familiar generalized Gamma distribution used in Ref. [9]. The reason for our choice is that we concentrate on the tail behavior of time distributions. According to our measurements the FPT, TTF and TTC distributions have fat tails, which can be well described by power laws. The generalized Gamma function has too slow convergence to a power law to describe the observed tails in the time range of our investigations. A detailed discussion of this problem is provided in Appendix B.

We also emphasize that in the present study we do not intend to discuss in detail the behavior on short time scales. We assume that this regime is simply characterized by the exponent λ' only to perform a quick and efficient fit. This choice will have no direct relevance to our main conclusions, which always apply to the tails of the distribution.

Nevertheless, in addition to the very good fit at large times the above formula gives for some cases an overall good description also at short times. Table II shows the results for all five stocks. We find that λ_{TTF} , which gives the asymptotic behavior of the distribution, ranges between 1.8 and 2.2 for up to $\Delta = 4$ ticks. This is greater than the value Ref. [8] found for NASDAQ. The exponent λ'_{TTF} varies between -0.4 and 0.6 . Finally T_{TTF} typically grows with Δ , as orders placed deeper into the book are executed later. We will return to this observation in Section IV C. For $\Delta > 4$ the small number of limit orders in our sample does not allow us to make reliable estimates for the shape of the distribution. Fig. 2 also gives a comparison of four further stocks (AZN, LLOY, SHEL and VOD) to show that our findings are quite general. The distribution of time to first fill is indistinguishable from time to fill.

For time to cancel one finds a similarly robust behavior, also shown in Fig. 2. Its distribution is again well fitted by the form

$$P_{\text{TTC};\Delta}(t) = \frac{C''t^{-\lambda_{\text{TTC}}}}{1 + [t/T_{\text{TTC}}(\Delta)]^{-\lambda_{\text{TTC}} + \lambda'_{\text{TTC}}}}, \quad (6)$$

where the long time asymptotics has an exponent λ_{TTC} ranging between 1.9 and 2.4. Unlike the case of λ_{TTF} , the

| stock | $\Delta = 1$ | | | $\Delta = 2$ | | | $\Delta = 3$ | | | $\Delta = 4$ | | |
|-------|--------------|------------|-----|--------------|------------|-----|--------------|------------|-----|--------------|------------|------|
| | λ | λ' | T |
| AZN | 1.50 | 0.14 | 58 | 1.50 | 0.22 | 140 | 1.50 | 0.18 | 240 | 1.49 | 0.11 | 350 |
| GSK | 1.52 | 0.16 | 62 | 1.52 | 0.18 | 230 | 1.50 | -0.02 | 390 | 1.48 | -0.21 | 520 |
| LLOY | 1.54 | 0.22 | 85 | 1.55 | 0.20 | 280 | 1.53 | 0.01 | 460 | 1.51 | -0.12 | 630 |
| SHEL | 1.52 | 0.20 | 83 | 1.53 | 0.27 | 160 | 1.51 | 0.02 | 360 | 1.51 | 0.00 | 450 |
| VOD | 1.57 | 0.43 | 150 | 1.54 | -0.19 | 450 | 1.49 | -0.69 | 720 | 1.51 | -0.66 | 1500 |

Table I: Parameters of the fitting function (3) for the distribution of first passage time for the five stocks. Δ is measured in ticks and all times are given in seconds. Typical standard errors for the quantities: ± 0.05 for λ_{FPT} , ± 0.05 for λ'_{FPT} , and $\pm 10\%$ for T_{FPT} .

measured values of λ_{TTC} are in agreement with those measured in Ref. [8] for NASDAQ. All results concerning the time to cancel are given in Table III.

As for the FPT, for both TTF and TTC the asymptotic power law behavior and the value of exponents is preserved if time is measured in transactions, see Appendix A.

B. Comparison of characteristic times

The emphasis of this paper is on the interplay between order execution, order cancellation and the first passage properties of price. To understand this relationship, consider the following argument proposed in Ref. [9]. Imagine that there are no cancellations. Let a buy order be placed at the price $b_0 - \Delta$, when the current best bid is at b_0 (the argument goes similarly for sell orders). How much time does it take until this order is executed? It is certain that the order cannot be executed before the best bid decreases to $b_0 - \Delta$, because until then there will always be more favorable offers in the book. On the other hand, once the price decreases to $b_0 - \Delta - \epsilon$ where ϵ is the tick size of the stock, it is certain, that all possible offers at the price $b_0 - \Delta$ have been exhausted, including ours. Therefore both time to fill and time to first fill for any order placed at a distance Δ from the best offer is greater than the first passage time of price to a distance Δ , and less than that to $\Delta + \epsilon$. Since this is true for every individual order, one expects the following inequality for the distribution functions of characteristic times:

$$\begin{aligned} \int_0^t P_{\text{FPT};\Delta}(t') dt' &\geq \int_0^t P_{\text{TTF};\Delta}(t') dt' \geq \\ \int_0^t P_{\text{TTFF};\Delta}(t') dt' &\geq \int_0^t P_{\text{FPT};\Delta+\epsilon}(t') dt'. \end{aligned} \quad (7)$$

Using the empirical distributions above, a straightforward calculation yields

$$\lambda_{\text{FPT}} = \lambda_{\text{TTFF}} = \lambda_{\text{TTF}}, \quad (8)$$

which is in clear disagreement with the data, where pronouncedly $\lambda_{\text{FPT}} < \lambda_{\text{TTF}} \approx \lambda_{\text{TTFF}}$. This inequality for the tail exponents means that one finds less orders with very long time to (first) fill than expected. The resolution of this apparent contradiction is that cancellations

have to be taken into account: Orders which would have to wait too long before being executed are often canceled and thus removed from the statistic. The measurement of the cancellation time distribution suffers from the same bias. The observed distribution of time to cancel does not characterize how traders would actually cancel their orders, because here the executed orders are missing from the statistics.

In Section V we will present a simple model that gives insight into the features pointed out so far. However, before doing so, we would like to present one further point concerning the empirical data.

C. The role of entry depth

How do order execution times change as a function of the entry depth Δ ? Similarly to first passage times, the empirical distributions found for time to fill/cancel have a slowly decaying tail such that the means might diverge. Therefore, in the following we will use the medians of all quantities as a measure of their typical value.

In Fig. 3 we show that the median of time to fill is empirically well described by

$$M[\text{TTF}] \propto \Delta^{1.4}, \quad (9)$$

which is quite different from the $M[\text{TTF}] \propto \Delta^2$ expected naively from Eq. (4) and a Brownian motion assumption, and also from the $\Delta^{1.5-1.8}$ behavior observed for the first passage time. We show that cancellations play an important role in these discrepancies.

Let us make a surrogate experiment with the data of the stock GSK. We select all filled orders, and from the time of their placement we calculate the *first* time when the transaction price becomes equal to or better than the price of the order. If one plots the median of this quantity versus the Δ of the orders, the resulting curve is indistinguishable from the median of time to fill [Fig. 3(left), curve labeled as "TTF/FPT filled ord"]. Thus the exponent 1.4 does not come from the difference between order executions and first passage times.

In another surrogate experiment we keep the time of order placements, but shuffle the Δ values between orders. This way we destroy correlations between volatility and order placement. We record the corresponding first passage times. The resulting curve is labeled as "FPT

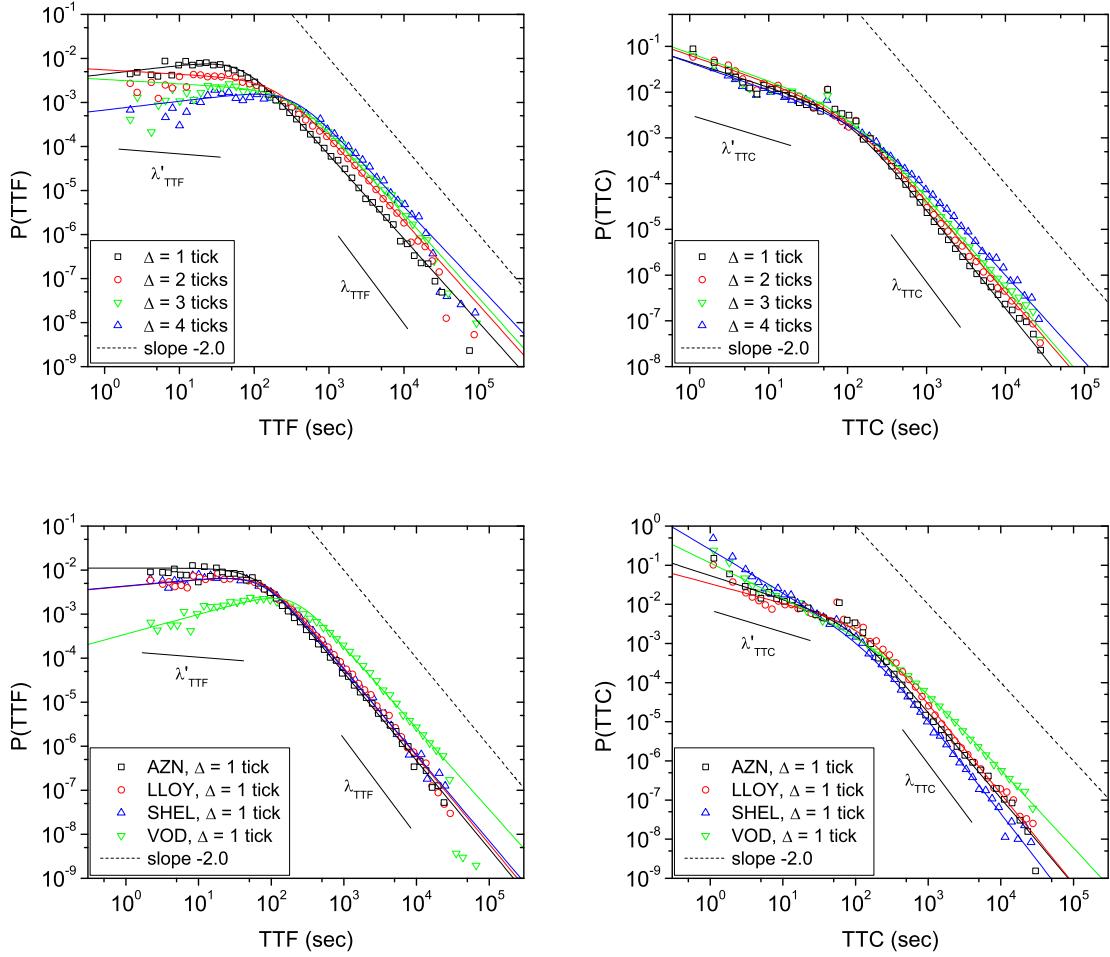


Figure 2: *Top left*: distribution of time to fill of GSK for $\Delta = 1 \dots 4$ ticks, and fits with Eq. (5). The dashed line is a power law with exponent -2.0 . *Top right*: distribution of time to cancel of GSK for $\Delta = 1 \dots 4$ ticks, and fits with Eq. (6). The dashed line is a power law with exponent -2.0 . *Bottom left*: comparison of distributions of time to fill for three typical stocks, $\Delta = 1$ tick, and fits with Eq. (5). The dashed line is a power law with exponent -2.0 . *Bottom right*: comparison of distributions of time to cancel for three typical stocks, $\Delta = 1$ tick, and fits with Eq. (6). The dashed line is a power law with exponent -2.0 .

| stock | $\Delta = 1$ | | | $\Delta = 2$ | | | $\Delta = 3$ | | | $\Delta = 4$ | | |
|-------|--------------|------------|-----|--------------|------------|-----|--------------|------------|-----|--------------|------------|-----|
| | λ | λ' | T |
| AZN | 2.0 | -0.0 | 65 | 1.9 | 0.0 | 100 | 1.8 | -0.0 | 120 | 1.9 | 0.0 | 200 |
| GSK | 1.9 | -0.2 | 68 | 1.9 | -0.2 | 150 | 1.8 | -0.4 | 190 | 1.8 | -0.3 | 320 |
| LLOY | 2.0 | -0.1 | 85 | 1.9 | -0.1 | 160 | 1.9 | -0.2 | 240 | 1.9 | -0.2 | 350 |
| SHEL | 1.9 | -0.1 | 77 | 1.9 | -0.2 | 110 | 1.9 | 0.0 | 270 | 1.8 | -0.1 | 250 |
| VOD | 1.8 | -0.4 | 190 | 1.8 | -0.5 | 490 | 1.8 | -0.4 | 980 | — | — | — |

Table II: Parameters of the fitting function (5) for the distribution of time to fill for the five stocks and $\Delta > 0$ ticks. All times are given in seconds. Data are missing where the statistics was inadequate for fitting. Typical standard errors for the quantities: ± 0.1 for λ_{TTF} , ± 0.1 for λ'_{TTF} , and $\pm 10\%$ for T_{TTF} .

shuff. all ord". This new curve now agrees with the first passage time of price [curve "FPT, price (book only)"] when $\Delta > 8$ ticks, which corresponds to a median time of about 1 – 2 hours. The origin of the anomalous Δ -dependence is, at least in the large Δ case, therefore the presence of cancellations. The explanation of other con-

tributions requires more involved arguments which are beyond the scope of this paper.

The dependence of median time to cancel on the entry depth Δ has a less clear functional form, as shown by Fig. 4. While $M[\text{TTC}]$ appears to be a monotonically increasing function of Δ , the curves for the different

| stock | $\Delta = 1$ | | | $\Delta = 2$ | | | $\Delta = 3$ | | | $\Delta = 4$ | | |
|-------|--------------|------------|-----|--------------|------------|------|--------------|------------|------|--------------|------------|------|
| | λ | λ' | T | λ | λ' | T | λ | λ' | T | λ | λ' | T |
| AZN | 2.2 | 0.6 | 87 | 2.2 | 0.6 | 90 | 2.2 | 0.6 | 85 | 2.2 | 0.7 | 100 |
| GSK | 2.2 | 0.5 | 110 | 2.0 | 0.5 | 90 | 1.9 | 0.5 | 94 | 1.9 | 0.6 | 170 |
| LLOY | 2.3 | 0.5 | 130 | 2.2 | 0.4 | 140 | 2.0 | 0.4 | 120 | 2.0 | 0.5 | 250 |
| SHEL | 2.4 | 1.1 | 150 | 2.3 | 1.1 | 140 | 2.3 | 1.1 | 68 | 2.2 | 1.0 | 56 |
| VOD | 2.0 | 0.9 | 300 | 2.1 | 0.8 | 1000 | 2.2 | 0.6 | 1500 | 1.9 | 0.5 | 1000 |

Table III: Parameters of the fitting function (6) for the distribution of time to cancel for the five stocks and $\Delta > 0$ ticks. All times are given in seconds. Data are missing when there were no orders at all, or the statistics was inadequate for fitting. Typical standard errors for the quantities: ± 0.1 for λ_{TTC} , ± 0.1 for λ'_{TTC} , and $\pm 25\%$ for T_{TTC} .

stocks show only a qualitative similarity. One of the reasons may be that different cancellation mechanisms are treated together.

V. A SIMPLE MODEL OF THE CHARACTERISTIC TIMES

The problem of the interplay between time to fill and time to cancel is an example for competing risks [16, 39]. In this framework mutually exclusive events are considered in time [16, 30]: in our case after its placement a limit order is either executed or canceled. Each of these events has its own probability distribution for the time when it will occur, but only the earliest one of the events is observed. In this section we present a simple joint model⁷ of limit order placement and cancellation that is of this type. We will see that the model gives predictions that can be tested against real data. Moreover, it also gives indications on the statistical properties of a quantity that is directly unobservable: the "lifetime" an agent is willing to wait for a limit order to be executed.

We make the following assumptions:

1. We consider one "representative agent" [41]. At time $t = 0$ the agent places a single buy⁸ limit order at a $\Delta > 0$ distance from the current best offer. (A generalization to $\Delta \leq 0$ is given in Section VII.) We treat all the other market participants on an aggregate level.
2. The agent is not willing to wait indefinitely for the order to be executed. Instead, at the time of placement the agent also decides about a cancellation (or more appropriately expiration) time t' for the order. This is a value drawn randomly from the distribution $P_{\text{LT};\Delta}(t')$. We will call this function the *lifetime distribution*. If the order is not executed until t' , then the order is canceled. The agent has

no additional cancellation strategy. This assumption is very restrictive (cf. Ref. [14]), but as Section VIII A will show, it does not affect our results significantly.

3. The market is very liquid and tick sizes are small. As a consequence,

- (a) before its execution, the effect of the agent's limit order on the evolution of the market price is negligible. This point neglects that traders reveal private information about their valuation of the stock by placing limit orders.
- (b) the interval between the time when the best bid reaches the order price and when the agent's order is executed is negligible. We also assume that such immediate execution is independent of the volume of the agent's order. A simple way to motivate that the volume present at a given price does not strongly affect execution times is to measure the typical ratio between time to fill and time to first fill as a function of the volume of the order. For at least 75% of the orders of any volume this is close to 1. The only exceptions can be very large orders with $\Delta = 1$. Here the price reaches the order quickly, but it takes about 20% longer to execute it completely (see also Ref. [9]). Moreover, for real limit orders the median time to fill does not depend too strongly on the volume of the order, except for very large volumes, see Fig. 5.

In our study we included SHEL and VOD which are known to have large tick/price ratios, so Assumption 3 would be invalid. Contrary to our expectations, we did not find any indication of anomalies like in other studies [14, 40, 42, 43], and the model proved useful for these stocks as well.

Under our assumptions one can write a joint density function that describes both the price diffusion process and cancellations. The probability $P_{\Delta}(t, t')$ that the price reaches an order placed at a distance $\Delta > 0$ from the current best offer at time t (and then it can be executed immediately), and that the agent decides to cancel the

⁷ Ref. [40] shows that similar arguments give a very good approximation for the average shape of the order book.

⁸ Note that throughout the paper we use the language of buy orders, but analogous definitions can be given for sell orders. All measurements include both buy and sell orders.

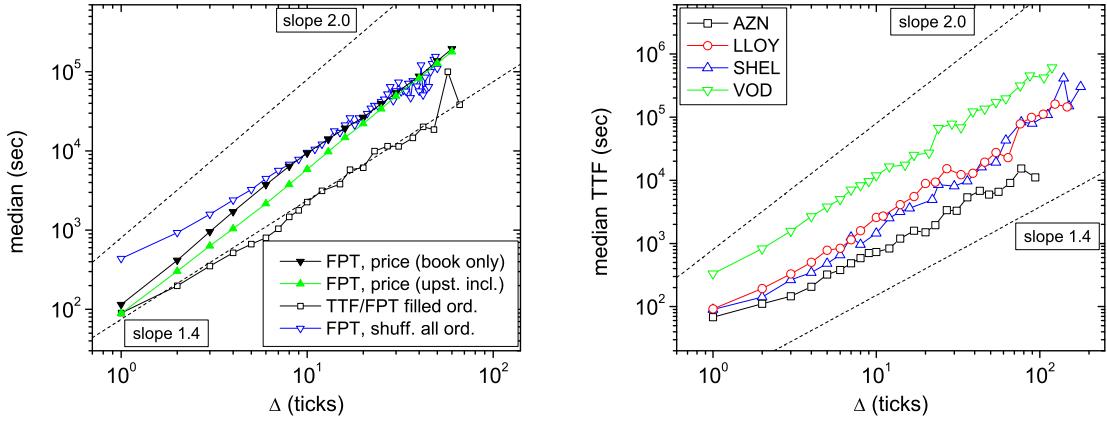


Figure 3: *Left*: The dependence of characteristic times on the entry depth Δ for GSK. Empty squares show time to fill, the curve for the surrogate first passage time assigned to filled orders is indistinguishable from this one. The dotted lines are $M[\cdot] \propto \Delta^{1.4}$ and $M[\cdot] \propto \Delta^{1.9}$. *Right*: The dependence of median time to fill on the entry depth Δ for four stocks. The curves show a strong similarity and they are significantly better fitted by $M[\text{TTF}] \propto \Delta^{1.4}$ than by the naive ansatz $M[\text{TTF}] \propto \Delta^2$ from Brownian motion.

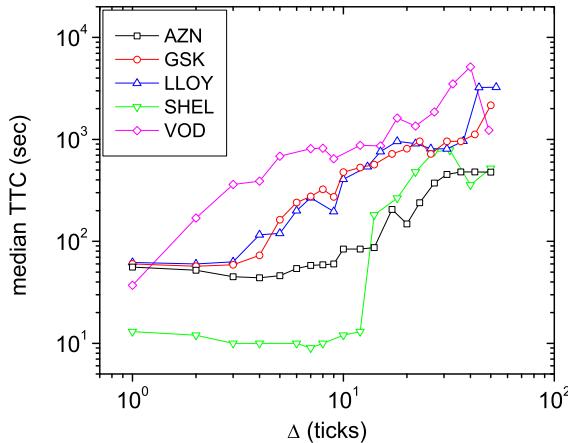


Figure 4: The dependence of median time to cancel on the entry depth Δ of the limit order. The curves have an increasing tendency and they are qualitatively similar across stocks. However, they do not follow any obvious functional form.

order a time t' can be written as a product of two independent distributions:

$$P_\Delta(t, t') = P_{\text{FPT};\Delta}(t)P_{\text{LT};\Delta}(t'). \quad (10)$$

For each limit order values of t and t' are drawn from P . The limit order is executed if $t < t'$ or it is canceled if $t > t'$. The two cases are illustrated in detail in Fig. 6.

VI. THE PREDICTIONS OF THE MODEL

Competing risk models are often estimated by the procedure introduced by Kaplan and Meier [44]. This is a statistically consistent, non-parametric method to estimate the marginal distributions P_{FPT} and P_{LT} from P_{TTF} and P_{TTC} under the assumption that execution and cancellation are independent as we already assumed in writing Eq. (10). We will now calculate these estimates in another, but strictly equivalent analytical way.

Let us denote distribution functions as follows:

$$P_{X;\Delta}(> t) = \int_t^\infty P_{X;\Delta}(\tau) d\tau, \quad (11)$$

where X can be any process introduced above (FPT, LT, TTF, TTFF, TTC). We will omit the lower index Δ for brevity. Let us first express the previously introduced quantities in terms of the joint probability $P_\Delta(t, t')$ and via Eq. (10). For executed orders $t < t'$, thus the distribution of time to fill is given by

$$P_{\text{TTF}}(t) = \frac{P_{\text{FPT}}(t)P_{\text{LT}}(> t)}{\int_0^\infty P_{\text{FPT}}(\tau)P_{\text{LT}}(> \tau) d\tau} = \mathcal{N}[P_{\text{FPT}}(t)P_{\text{LT}}(> t)]. \quad (12)$$

We introduced the operator $\mathcal{N}[\cdot]$, which normalizes a function to an integral of 1. Symmetrically for time to cancel $t < t'$:

$$P_{\text{TTC}}(t) = \frac{P_{\text{FPT}}(> t)P_{\text{LT}}(t)}{\int_0^\infty P_{\text{FPT}}(> \tau)P_{\text{LT}}(\tau) d\tau} = \mathcal{N}[P_{\text{FPT}}(> t)P_{\text{LT}}(t)]. \quad (13)$$

As (12) and (13) are two equations with only one unknown function, namely the lifetime distribution $P_{\text{LT}}(t)$,

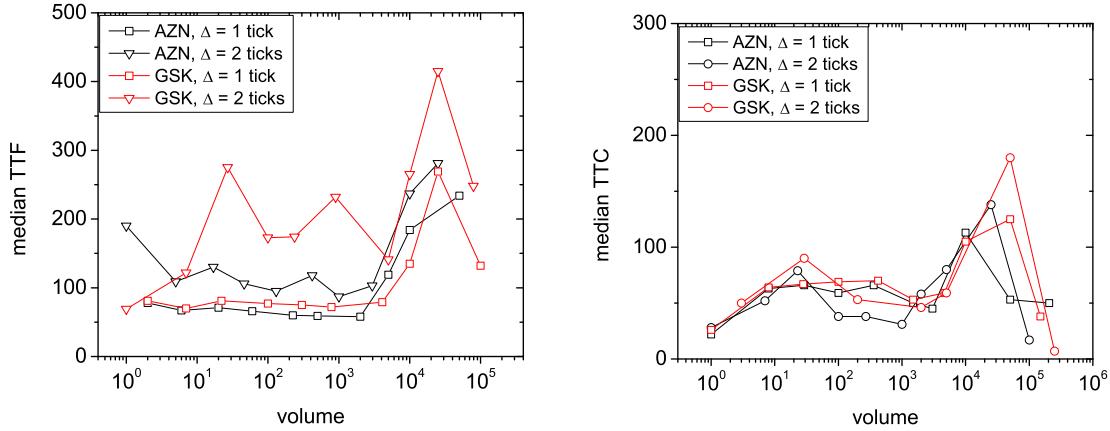


Figure 5: *Left*: median time to fill for AZN and GSK, $\Delta = 1$ and 2 ticks, as a function of order volume. The value does not depend strongly on order size except for very large orders. *Right*: median time to cancel for AZN and GSK, $\Delta = 1$ and 2 ticks, as a function of order volume. The value does not depend strongly on order size.

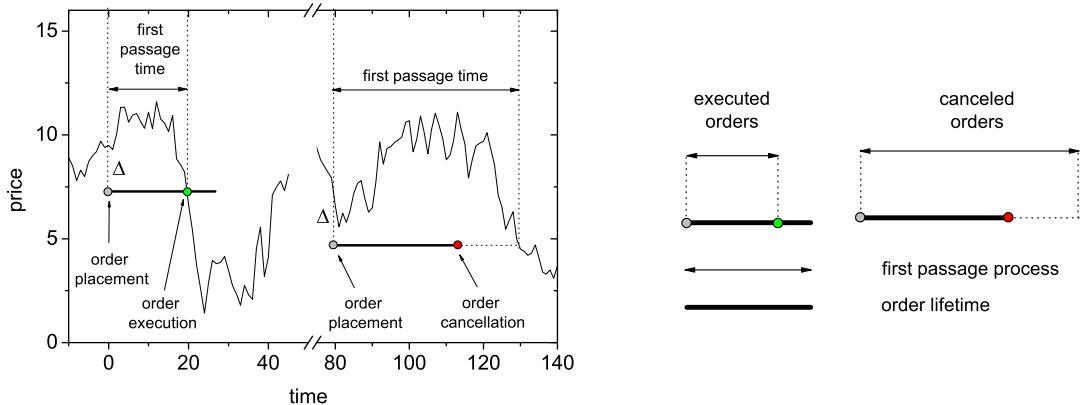


Figure 6: *Left*: The scheme of the diffusive model for execution times (example with buy orders). Orders are indicated by thick horizontal lines. The order is placed at a distance Δ below the current best bid. At the time of its placement the order is assigned a lifetime (the length of the thick line). If the bid crosses the line, then the order is executed at the time of crossing. The time between order placement and the crossing is the first passage time of the bid price to a distance Δ . If there is no crossing, the order is canceled at its cancellation time (the end of the thick line). *Right*: Orders are executed when the first passage time is less, and canceled when larger than the intended lifetime.

one can calculate that from, e.g., Eq. (12), and then see if the solution is consistent with Eq. (13). We can express from Eq. (12), that

$$P_{LT}(>t) \propto \frac{P_{TTF}(t)}{P_{FPT}(t)} \quad (14)$$

and thus

$$P_{LT}(t) = -\frac{d}{dt}P_{LT}(>t) = -\mathcal{N} \left[\frac{d}{dt} \frac{P_{TTF}(t)}{P_{FPT}(t)} \right]. \quad (15)$$

It is also possible to estimate the same quantity directly

from Eq. (13):

$$P_{LT}(t) = \mathcal{N} \left[\frac{P_{TTC}(t)}{P_{FPT}(>t)} \right]. \quad (16)$$

Let us eliminate the lifetime distribution, and substitute the large t asymptotic power law behavior of all probabilities. After simple calculations one finds that

$$\lambda_{TTF} = \lambda_{TTC}. \quad (17)$$

Then we substitute this result back into Eq. (14) to find that the lifetime distribution also has to decay asymptotically as a power law:

$$P_{LT}(t) \propto t^{-\lambda_{LT}}, \quad (18)$$

with

$$\lambda_{LT} = \lambda_{TTF} - \lambda_{FPT} + 1 = \lambda_{TTC} - \lambda_{FPT} + 1. \quad (19)$$

Eq. (17) is in good agreement with the results of Section IV, where $\lambda_{TTF} = 1.8 - 2.2$, and $\lambda_{TTC} = 1.9 - 2.4$. This is a clear improvement compared to Eq. (8). The introduction of the simplest possible cancellation model gives a good prediction for the difference between the exponents describing the asymptotics of the first passage time and time to fill.

Moreover, one can now observe the hidden distribution of lifetimes. By substituting the typical values into Eq. (19), one gets $\lambda_{LT} \approx 1.6$. In comparison, a paper by Bordan and Bouchaud [45] describes a GARCH-like model obtained by introducing a distribution of traders' investment horizons and the model reproduces empirical values of volatility correlations for $\lambda_{LT} = 1.15$, which is not far from our estimate. More recently it has been shown [46] that the limit order price probability distribution is consistent with the solution of an utility maximization problem in which the limit order lifetime is power law distributed with an exponent $\lambda_{LT} \simeq 1.75$. The origin of the power law distribution of limit order lifetimes is not clear. Unfortunately the data do not allow us to separate individual traders. Therefore we do not know whether such a result arises from the broad distribution of the time horizons of each trader, or simply a distribution of traders with different investment strategies. Based on an empirical investigation at the broker level, in Ref. [46] it is argued that heterogeneity of investors could be the determinant of the power law lifetime distribution. Notice, however, two points: (i) We are not speaking about how long the investors *hold* the stock. Instead, P_{LT} is the distribution of *how long investors are willing to wait* for their limit orders to be executed and before they cancel or revise their offers. (ii) None of the limit orders we are discussing here are truly long-term. Even the orders with relatively long lifetime spend at most a few days in the book.

VII. AN EXTENSION TO $\Delta \leq 0$

So far we only considered orders with prices which were worse than the best offer at the time of their placement, i.e., $\Delta > 0$. However, this group only accounts for less than half of the actual limit orders. Measurements for $\Delta \leq 0$ orders give the surprising result that these execution times are described by statistics very similar to those for $\Delta > 0$. One example stock (GSK) is shown in Fig. 7(left). The results of our fitting procedure performed with Eq. (5) are given in Table IV for all five stocks.

According to our model, these orders should have been executed within a negligible time of their placement. While this is true for a number of them, certainly not for all. Let us assume that we are placing a new buy limit order. If our order has $\Delta = 0$, then it will be among the best offers at the time of its placement. If our order has

$\Delta < 0$, then it becomes the single best offer in the book, and hence it will trade with certainty if the next event is a buy market order. Why can our order still take a long time before being executed? The answer is naturally that before our order is executed, a new buy limit order may enter the book. If this new order has $\Delta < 0$ (where Δ now has to be measured from our order), it means that it has an even better price than our order and it will gain priority of execution. On the other hand, our order now effectively has $\Delta > 0$, and the original model can be applied.

In order to test such a hypothesis, we carried out the following calculation. For the sake of simplicity, we will consider the time to first fill instead of time to fill. Section V argued that for the majority of orders the difference between the two is negligible. From the time of its placement, we tracked every single at least partially filled $\Delta \leq 0$ order until the time it was first filled. We defined the reduced entry depth (Δ') and the reduced time to first fill ($TTFF'$) for these orders as follows

1. For orders, where from their placement to their first fill there were no even more favorable orders both placed and then at least partially filled, $\Delta' = 0$ and $TTFF' = TTFF$.
2. For orders where after their placement but before their first fill there was at least one new, more favorable order introduced with $\Delta_{new} < 0$ and then this new order was at least partially filled, we selected the first of such new orders placed after the original one and set $\Delta' = -\Delta_{new}$. Thus, Δ' is the new position of the original order, after the new one was placed. $TTFF'$ is defined as the time to first fill of our order measured from the placement of this new order.

The typical distribution of $TTFF'$ for different groups in Δ' is shown in Fig. 7(right). For orders with $\Delta' = 0$ this is – except for here uninteresting very short times – well described by a stretched exponential distribution $P_{TTFF'}(t) = \frac{1}{25} \exp\left[-\left(\frac{t}{6}\right)^{1/2}\right]$. These are the orders, where there was no better offer made, and hence their execution times were purely determined by the incoming market orders. The distribution is very close to the distribution of the times between two consecutive transactions of the stock [see Fig. 7(right)].

For orders with $\Delta' > 0$, one recovers the results of the previous sections, and the distribution of reduced time to first fill asymptotically decays as a power law with a power close to 2.0. Eqs. (17) and (19) are expected to be valid for orders with $\Delta < 0$ and $\Delta' > 0$ as well, given that we use them in terms of Δ' and $TTFF'$.

As a summary, time to first fill for orders with $\Delta \leq 0$ is a two-component process. If there is no better order placed before the first fill, then time to first fill is basically identical to the waiting time distribution between opposite market orders. If there is a better offer submitted, then the order effectively becomes $\Delta > 0$, and the

| stock | $\Delta = 0$ | | | $\Delta = -1$ | | | $\Delta = -2$ | | |
|-------|--------------|------------|-----|---------------|------------|-----|---------------|------------|-----|
| | λ | λ' | T | λ | λ' | T | λ | λ' | T |
| AZN | 2.2 | 0.6 | 110 | 2.2 | 1.0 | 230 | 2.1 | 1.2 | 320 |
| GSK | 2.2 | 0.5 | 110 | 2.1 | 1.1 | 180 | 2.2 | 1.3 | 410 |
| LLOY | 2.2 | 0.5 | 120 | 2.1 | 1.0 | 150 | 2.0 | 1.2 | 220 |
| SHEL | 2.2 | 0.4 | 110 | 2.1 | 1.0 | 120 | 2.1 | 1.0 | 140 |
| VOD | 2.1 | 0.5 | 160 | 2.0 | 1.1 | 130 | — | — | — |

Table IV: Parameters of the fitting function (5) for the distribution of time to fill for the five stocks and $\Delta \leq 0$ ticks. All times are given in seconds. Data are missing when there were the statistics was inadequate for fitting. Typical standard errors for the quantities: ± 0.1 for λ_{TTF} , ± 0.1 for λ'_{TTF} , and $\pm 25\%$ for T_{TTF} .

| stock | $\Delta = 0$ | | | $\Delta = -1$ | | | $\Delta = -2$ | | |
|-------|--------------|------------|-----|---------------|------------|-----|---------------|------------|-----|
| | λ | λ' | T | λ | λ' | T | λ | λ' | T |
| AZN | 2.3 | 0.6 | 130 | 2.0 | 0.7 | 90 | 1.9 | 0.8 | 72 |
| GSK | 2.1 | 0.6 | 130 | 1.9 | 0.7 | 90 | 1.8 | 0.8 | 50 |
| LLOY | 2.2 | 0.5 | 120 | 1.9 | 0.7 | 70 | 1.8 | 0.9 | 85 |
| SHEL | 2.3 | 1.0 | 220 | 2.2 | 1.0 | 130 | 2.0 | 1.0 | 160 |
| VOD | 2.0 | 0.7 | 200 | 1.8 | 0.8 | 120 | — | — | — |

Table V: Parameters of the fitting function (6) for the distribution of time to cancel for the five stocks and $\Delta \leq 0$ ticks. All times are given in seconds. Data are missing when there were no orders at all, or the statistics was inadequate for fitting. Typical standard errors for the quantities: ± 0.1 for λ_{TTC} , ± 0.1 for λ'_{TTC} , and $\pm 25\%$ for T_{TTC} .

diffusion approximation applies. As this latter process has a much fatter tail than the former one, long waiting times and the tail exponent of the joint process are again dominated by a first passage process.

VIII. DISCUSSION

A. Lifetime distribution

Before discussing the results let us analyze the most important simplifying assumption of our model, namely the independence of the lifetime of the order from the evolution of price. This would mean that traders decide about an expiry time of their limit orders at the time of their placement, and then do not cancel them earlier, which resembles the random cancellation process as introduced in Ref. [13]. In order to see the relevance of our assumption one should calculate the cross-correlation coefficient of first passage times and the lifetime process. However, as mentioned in Section V we are limited by the fact that the lifetime is hidden. It is not possible to calculate cross-correlations between time to fill and time to cancel either, because for the same order one cannot observe both variables. This issue is related to the identifiability problem of competing risks [30].

We suggest the following approach to resolve the above issue: Let us consider canceled orders only. There one can observe the values of the lifetime, because they were realized as an actual time to cancel. Moreover, our model assumed, that the order would have been executed at the first passage time (the time of the first transaction at the order's or a better price). Now it is possible to quantify cross-correlations between these two quantities, but one has to keep in mind three points. (Note that we will

consider orders with $\Delta = 1$ to have the largest possible sample.)

1. For very short times the price dynamics is dominated by bid-ask bounce, and other non-diffusive processes [47]. Our model is not valid in this regime, because rapid order executions are not governed by a first passage process. Hence we discard all orders which were canceled within $L = 4$ minutes of their placement.
2. In order to avoid problems arising from the possible non-existence of the moments of the distributions, we choose to evaluate Spearman's rank-correlation coefficient⁹ (ρ), instead of Pearson's correlation coefficient. The quantity ρ has further favorable statistical properties, for example it is not very sensitive to extreme events.
3. As we can only consider canceled orders, we know that $\text{FPT} > \text{LT}$. This constraint alone, and regardless of the choice of correlation measure, will cause strong positive correlations between the two quantities. Even if FPT and LT are independent, the conditional joint distribution reads

$$P(\text{FPT} = t, \text{LT} = t' | \text{FPT} > \text{LT}) = \mathcal{N}[\Theta(t - t') P_{\text{FPT}}(t) P_{\text{LT}}(t')], \quad (20)$$

⁹ This is defined by first, for both quantities separately, replacing each observation by its rank in the sample (i.e., assigning 1 to the largest observation of first passage time, 2 to the second largest, etc., and then repeating the procedure for lifetimes). Then the usual cross-correlation coefficient is calculated for the ranks [48].

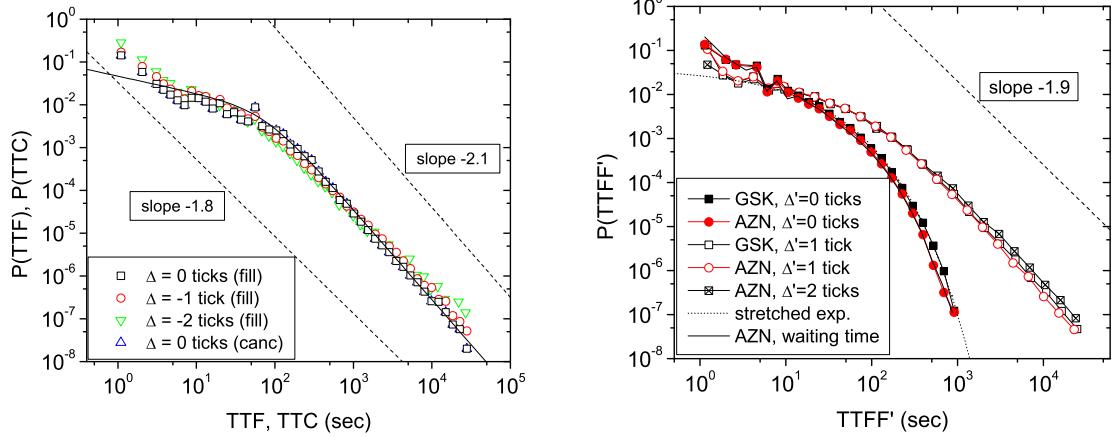


Figure 7: *Left:* Examples of the distribution of time to fill/cancel for $\Delta \leq 0$ buy limit orders of GSK. *Right:* The distribution of reduced time to first fill (TTFF') as a function of the reduced entry depth Δ' . For orders where $\Delta' = 0$, the tail of the distribution is well fitted by the stretched exponential $\frac{1}{25} \exp\left[-\left(\frac{t}{6}\right)^{1/2}\right]$. Where $\Delta' > 0$, the distribution decays asymptotically as a power law with an exponent close to -2.0 . The solid line is the distribution of waiting times between two consecutive trades of GSK.

where Θ is the Heaviside step function. Due to our restricted observations this is clearly not a product of two independent densities.

Instead, a more convenient null hypothesis is to measure the correlations between FPT/LT and LT. $L = 4$ min was chosen such that for $\Delta = 1$ the distribution of the first passage time is well described by the power law

$$P_{\text{FPT}}(t|t > L) \sim \frac{\lambda_{\text{FPT}} - 1}{L^{\lambda_{\text{FPT}} - 1}} t^{-\lambda_{\text{FPT}}}. \quad (21)$$

If FPT and LT are independent, then

$$\begin{aligned} P(\text{FPT}/\text{LT} = x, \text{LT} = t' | \text{FPT} > \text{LT}) &= \\ \mathcal{N}[\Theta(x-1)P_{\text{FPT}}(xt')P_{\text{LT}}(t')] &= \\ \mathcal{N}[\Theta(x-1)x^{-\lambda_{\text{FPT}}}] \times \mathcal{N}[P_{\text{FPT}}(t')P_{\text{LT}}(t')]. \end{aligned} \quad (22)$$

Eq. (21) was used for the second equality. The final result is a product form in functions of x and of t' , which means that FPT/LT is independent from LT, given that we restrict ourselves to $\text{FPT} > \text{LT}$. Remember that the only assumption for this result is that first passage times are asymptotically power law distributed, which seems to hold very well in our data down to $L \approx 4$ min.

We calculated Spearman's rank correlations between FPT/LT and LT in our restricted sample for various stocks, this we will denote by ρ_{res} . Results are summarized in Table VI. One finds negative correlation between the two quantities at all usual significance levels.¹⁰ This

means that those limit orders that would have been executed later were canceled earlier, i.e., that traders update their decision on when to cancel a limit order by tracking the price path. This is in line with the results of Ref. [14]. To prove that this value of ρ truly comes from correlations, we generated surrogate datasets by randomizing the pairs FPT/LT and LT while keeping the constraint $\text{FPT} > \text{LT}$. According to Table VI this completely destroys the correlations between FPT/LT and LT, $\rho_{\text{surr}} = 0$.

It is important to remember that this value of ρ_{res} is not the actual correlation coefficient between the first passage time and the lifetime process. To quantify the true value of cross-correlations, we introduce ρ_{true} which is Spearman's rank-correlation coefficient between LT and FPT. While this cannot be measured directly, there is a procedure to estimate it from a known value of ρ_{res} based on Monte Carlo simulation. Let us assume that FPT and LT are adequately described by power law distributions with the known tail exponents. We model the cross-correlation between the two processes by copulas (see Ref. [50]). Morgenstern's copula reads

$$P(>t, >t') = P_{\text{FPT}}(>t)P_{\text{LT}}(>t') \{1 + 3\rho_{\text{true}}[1 - P_{\text{FPT}}(>t)][1 - P_{\text{LT}}(>t')]\}, \quad (23)$$

with some $-1/3 < \rho_{\text{true}} < 1/3$, while Frank's copula

¹⁰ The error bars were estimated by the bootstrapping procedure

suggested in Ref. [49] (for more details see Refs. therein).

assumes

$$P(>t, >t') = \frac{1}{\alpha} \ln \left[1 + \frac{(e^{\alpha P_{\text{FPT}}(>t)} - 1)(e^{\alpha P_{\text{LT}}(>t')} - 1)}{e^\alpha - 1} \right], \quad (24)$$

with some $-\infty < \alpha < \infty$. Here $P(>t, >t') = \int_t^\infty d\tau \int_{t'}^\infty d\tau' P(\tau, \tau')$ which is the joint distribution function.

Monte Carlo measurements based on random pairs from these copulas suggest a nearly linear relationship between the true and the restricted correlation coefficients. With the substitution of the typical values of λ_{FPT} and λ_{LT} one finds that

$$\rho_{\text{true}} = r \times \rho_{\text{res}}, \quad (25)$$

where $r \approx 1.66$ for Morgenstern's and $r \approx 1.55$ for Frank's copula. The resulting estimates are given in Table VI. Naturally, the shuffled surrogate datasets yield $\rho_{\text{true}} = \rho_{\text{res}} = 0$.

These calculations have shown that there is a strong negative correlation between the first passage time and the lifetime of an order in agreement with Ref. [14] but contrary to our model assumption 2 and Eq. (10). So the key question is: How much does the presence of this correlation affect the predictions of our model? We performed a series of Monte Carlo simulations of the execution and cancellation processes by using the empirically observed value of tail exponents and cross correlations (Table VI). We found that for a fixed value of λ_{FPT} and λ_{LT} the introduction of such correlations increases the values of λ_{TTF} and λ_{TTC} by about 0.1, which is comparable to the error bars of our estimates, and the power law behavior is well preserved. Moreover, the central part of our arguments, Eq. (17), remains valid. Thus the presence of a dynamic cancellation strategy does not significantly affect the validity of our model.

B. Conclusions

In this paper we focused on the tails of the distributions of characteristic times in the limit order book. Our empirical observations, based on five highly liquid stocks on the London Stock Exchange, underline the importance of cancellations when comparing the first passage time to the time to execute an order. We found that the distributions follow asymptotically power laws for the first passage time, the time to (first) fill and time to cancel. The differences between the statistical properties of

these characteristic times are informative of the interdependence of order executions and cancellations. These observations are quite robust and can be seen as "stylized facts" characterizing the order book.

We did not find significant difference between the behavior of buy and sell orders, in contrast with Refs. [9, 51] for US markets, but in accord with Ref. [11] for the case of Ericsson stock traded at the Stockholm Stock Exchange. We are therefore not able to conclude whether the symmetric behavior we observe in the London Stock Exchange is common to most markets or specific to some of them or to certain time periods.

In addition to the empirical findings summarized in Tables I, II and III we introduced a model, where order execution times are related to the first passage time of price, and orders are canceled randomly with lifetimes that are asymptotically power law distributed. This can be considered as the simplest possible model to take cancellations into account. In this framework we showed that the characteristic exponents of the asymptotic power law behavior of the first passage time, the time to (first) fill and time to cancel are related to each other by simple rules which are in agreement with our empirical observations. These results are in contrast with another study (the NASDAQ data investigated in Ref. [8]). Therefore further investigations are needed to clarify whether or not our findings are market specific.

The observed heterogeneity of cancellation times may be driven by traders having different time horizons or by traders following different cancellation strategies in different market environments. Methods that can discriminate between these mechanisms represent a major objective for future research.

Acknowledgments

We would like to thank two anonymous referees for useful comments and suggestions. ZE is grateful to Jean-Philippe Bouchaud for discussions of the order book, to Michele Tumminello for advice on bootstrapping, and to Ingve Simonsen for help with the measurement of first passage times. The hospitality of l'Ecole de Physique des Houches and Capital Fund Management is also thankfully acknowledged. This work was supported by COST-STSM-P10-917 and OTKA T049238. FL and RNM acknowledge support from MIUR research project "Dinamica di altissima frequenza nei mercati finanziari" and NEST-DYSONET 12911 EU project.

[1] B. Biais, L. Glosten, and C. Spatt, *Journal of Financial Markets* **8**, 217 (2005).
[2] B. Biais, P. Hillion, and C. Spatt, *Journal of Finance* **50**, 1655 (1995).
[3] P. Handa and R. Schwartz, *Journal of Finance* **51**, 1835 (1996).
[4] L. Harris and J. Hasbrouck, *Journal of Financial and Quantitative Analysis* **31**, 213 (1996).
[5] K. Kavajecz, *Journal of Finance* **54**, 747 (1999).
[6] P. Sandas, *Review of Financial Studies* **14**, 704 (2001).

| stock | ρ_{res} | ρ_{surr} | ρ_{true} (Morg.) | ρ_{true} (Frank) | number of points |
|-------|---------------------|----------------------|------------------------------|------------------------------|------------------|
| AZN | -0.12 ± 0.02 | -0.001 ± 0.001 | -0.19 ± 0.03 | -0.18 ± 0.03 | 3277 |
| GSK | -0.13 ± 0.01 | 0.000 ± 0.001 | -0.21 ± 0.02 | -0.20 ± 0.02 | 8573 |
| LLOY | -0.10 ± 0.01 | 0.000 ± 0.001 | -0.16 ± 0.02 | -0.15 ± 0.02 | 8201 |
| SHEL | -0.13 ± 0.02 | -0.001 ± 0.001 | -0.21 ± 0.03 | -0.19 ± 0.03 | 2791 |
| VOD | -0.134 ± 0.008 | 0.000 ± 0.001 | -0.22 ± 0.01 | -0.21 ± 0.01 | 16392 |

Table VI: Estimates of the correlation between first passage times and the lifetimes for orders with entry depth $\Delta = 1$ tick. We discarded orders executed or canceled within $L = 4$ minutes of their placement. The errors represent the standard deviation estimated from 300 bootstrap samples.

- [7] S. Maslov and M. Mills, *Physica A* **299**, 234 (2001).
- [8] D. Challet and R. Stinchcombe, *Physica A* **300**, 285 (2001).
- [9] A. Lo, A. MacKinlay, and J. Zhang, *Journal of Financial Econometrics* **65**, 31 (2002).
- [10] M. Potters and J.-P. Bouchaud, *Physica A* **324**, 133 (2003).
- [11] B. Hollifield, R. A. Miller, and P. Sandas, *Review of Economic Studies* **71**, 1027 (2004).
- [12] J. D. Farmer, L. Gillemot, F. Lillo, S. Mike, and A. Sen, *Quantitative Finance* **4**, 383 (2004).
- [13] J. D. Farmer, P. Patelli, and I. I. Zovko., *Proc. Natl. Acad. Sci. USA* **102**, 2254 (2005).
- [14] S. Mike and J. D. Farmer, *Journal of Economic Dynamics and Control* **32**, 200 (2008).
- [15] P. Weber and B. Rosenow, *Quantitative Finance* **5**, 357 (2005).
- [16] B. Hollifield, R. A. Miller, P. Sandas, and J. Slive, *Journal of Finance* **61**, 2753 (2006).
- [17] A. Ponzi, F. Lillo, and R. N. Mantegna, *physics/0608032* (2006).
- [18] C. Parlour, *Review of Financial Studies* **11**, 789 (1998).
- [19] M. G. Daniels, J. D. Farmer, G. Iori, and E. Smith, *Phys. Rev. Lett.* **90**, 108102 (2003).
- [20] T. Foucault, *Journal of Financial Markets* **2**, 99 (1999).
- [21] T. Foucault, O. Kadan, and E. Kandel, *Review of Financial Studies* **18**, 1171 (2005).
- [22] I. Rosu, A dynamic model of the limit order book, Working paper, University of Chicago.
- [23] L. Glosten, *Journal of Finance* **49**, 1127 (1994).
- [24] S. Chakravarty and C. Holden, *Journal of Financial Intermediation* **4**, 213 (1995).
- [25] D. Seppi, *Review of Financial Studies* **10**, 704 (1997).
- [26] H. Luckock, *Quantitative Finance* **3**, 385 (2003).
- [27] C. Chiarella and G. Iori, *Quantitative Finance* pp. 346–353 (2002).
- [28] M. L. Calzi and P. Pellizzari, *Quantitative Finance* **3**, 470 (2003).
- [29] B. Chakrabarty, Z. Han, K. Tyurin, and X. Zheng (2006), CAEPR Working Paper 2006-015.
- [30] T. Bedford, *Competing risk modeling in reliability* (World Scientific Publishing, 2005), vol. 10 of *Modern Statistical and Mathematical Methods in Reliability: Series on Quality, Reliability and Engineering Statistics*.
- [31] W. Feller, *An introduction to probability theory*, vol. 1 (Wiley & Sons, 1967), 3rd ed.
- [32] A. V. Chechkin, R. Metzler, V. Y. Gonchar, J. Klafter, and L. V. Tanatarov, *J. Phys. A: Math. Gen.* **36**, L537 (2003).
- [33] S. Redner, *A Guide to First-Passage Processes* (Cambridge University Press, 2001).
- [34] I. M. Sokolov and R. Metzler, *J. Phys. A: Math. Gen.* **37**, L609 (2004).
- [35] E. Scalas, *Physica A* **362**, 225 (2005).
- [36] M. Montero, J. Perelló, J. Masoliver, F. Lillo, S. Micciché, and R. N. Mantegna, *Phys. Rev. E* **72**, 056101 (2005).
- [37] I. Simonsen, M. Jensen, and A. Johansen, *Eur. Phys. J. B* **27**, 583 (2002).
- [38] V. Seshadri and B. J. West, *PNAS* **79**, 4501 (1982).
- [39] D. Bernoulli, *Essai d'une nouvelle analyse de la mortalité causee par la petite verole* (Mem. Math. Phy. Acad. Royal Sciences, Paris, 1766).
- [40] M. Wyart, J.-P. Bouchaud, J. Kockelkoren, M. Potters, and M. Vettorazzo, *Quantitative Finance* **8**, 41 (2007).
- [41] A. P. Kirman, *Journal of Economic Perspectives* **6**, 117 (1992).
- [42] Z. Eisler and J. Kertész, *Eur. Phys. J. B* **51**, 145 (2006), (also *physics/0508156*).
- [43] Z. Eisler and J. Kertész, *Europhys. Lett.* **77**, 28001 (2007).
- [44] E. L. Kaplan and P. Meier, *Journal of the American Statistical Association* **53**, 457 (1958).
- [45] L. Borland and J.-P. Bouchaud, *physics/0507073* (2005).
- [46] F. Lillo, *Eur. Phys. J. B* **55**, 453 (2007).
- [47] R. Roll, *Journal of Finance* **39**, 1127 (1984).
- [48] J. C. Lee, *Statistics for business and financial economics* (World Scientific Publishing, 2000), 2nd ed.
- [49] F. Schmid and R. Schmidt, *Proc. Compstat.* **06**, 759 (2006).
- [50] E. W. Frees and E. A. Valdez, *North American Actuarial Journal* **2**, 1 (1998).
- [51] J.-W. Cho and E. Nelling, *Financial Analysts Journal* **56**, 28 (2000).
- [52] B. Hill, *Annals of Statistics* **3**, 1163 (1975).
- [53] P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modeling Extremal Events for Insurance and Finance* (Springer-Verlag, Berlin, 1997).

Appendix A: RESULTS IN TRANSACTION TIME

The typical time between transactions strongly depends on market conditions and it is very far from strictly stationary. This fact, also closely related to volatility clustering, could influence the distribution of first passage times, time to fill and time to cancel. Many recent studies measure time in transactions in order to remove fluctuations in trading activity. In order to better understand the role of activity fluctuations, we repeated our calculations in transaction time, but we did not find any changes that affect the conclusions of our paper. Fig. 8 shows comparisons between real time and transaction time for

the probability distributions of FPT, TTF and TTC (for the stock GSK, $\Delta = 1$ tick). The short time regime is quite different, while for long times the fluctuations in trading activity are less relevant, and all the distributions remain power laws asymptotically. The changes in the values of the tail exponents are also small. The bottom right panel of Fig. 8 compares P_{FPT} , P_{TTF} and P_{TTC} in transaction time. Our arguments still hold, as $\lambda_{\text{FPT}} < \lambda_{\text{TTF}} \approx \lambda_{\text{TTC}}$.

Appendix B: FITTING FUNCTIONS FOR THE DISTRIBUTIONS OF TIME TO FILL AND TIME TO CANCEL

In this Appendix we present a critical discussion regarding our decision to fit the empirical distributions of FPT, TTF and TTC with the functional form presented in Eq. (5). In our preliminary investigations, we fitted the distribution of FPT, TTF and TTC with two different distributions. The first one was the one we consider throughout the paper, i.e.,

$$P_Z(t) = \frac{C't^{-\lambda}}{1 + (t/T)^{-\lambda+\lambda'}}. \quad (\text{B1})$$

The second one was the generalized Gamma distribution

$$P_G(t) = \frac{\lambda|p|\kappa^\kappa(\lambda t)^{p\kappa-1} \exp[-(\lambda t)^p\kappa]}{\Gamma(\kappa)}, \quad (\text{B2})$$

which has been used in some of the existing studies on TTF (e.g. Ref. [9]). Another common form, the Weibull distribution, is a special case of Eq. (B2) for $\kappa = 1$. Our empirical analysis shows that the Weibull distribution fits the data poorly and it will not be considered in this Appendix. For large values of t the density of Eq. (B1) behaves as

$$P_Z(t) \sim \frac{1}{t^\lambda} \quad (\text{B3})$$

The asymptotic behavior of $P_G(t)$ depends on the sign of the parameter p . If $p < 0$ (as for the investigated data) it can be written as

$$P_G(t) \sim \frac{\exp(-c/t^{|p|})}{t^{1+|p|\kappa}} \quad (\text{B4})$$

where c is a constant. Thus the generalized Gamma distribution, similarly to Eq. (B1), is consistent with a power law tail, although it is modulated by an exponential function which becomes less and less important as $t \rightarrow \infty$. In order to estimate the optimal parameters of the distributions we used a Maximum Likelihood Estimator (MLE). For illustrative purposes, here we consider

the case of TTF for AZN and $\Delta = 0$ but the results are similar for other stocks, other values of Δ and for both TTF and TTC.

Fig. 9(left) shows the distribution of TTF for AZN and $\Delta = 0$ together with fits by Eqs. (B1) and (B2). Both P_Z and P_G give a good fit both in the tail and in the body of the distribution. One finds that P_G has a slightly larger likelihood \mathcal{L} than P_Z . Since the two distributions have the same number of parameters (degrees of freedom) the likelihoods can be compared directly. However if one computes the tail exponents of the distribution from the fitted parameters one finds a puzzling result. The tail exponent obtained from the generalized Gamma distribution fit is 4.5, whereas the tail exponent obtained from the P_Z fit is 2.2. Such a difference in the exponent should be detectable in data. Still, Fig. 9(left) shows that both distributions fit the tail reasonably well. The reason of this contradiction is shown in the inset of Fig. 9(left). This plots the local tail exponent of the generalized Gamma distribution, given by $d[\log P_G(t)]/d[\log t]$, as a function of t . The local exponent of the generalized Gamma distribution converges extremely slowly to the asymptotic value 4.5 and in the range of the TTF from 10^3 to 10^4 the local exponent is between 2 and 3, which is approximately consistent with the values obtained from P_Z .

As we have repeatedly stated, in this paper we are interested in the tail behavior of the distribution of the time to fill and time to cancel. The analysis summarized in Fig. 9(left) shows that the parameters estimated from a fit to a generalized Gamma distribution are not suitable to estimate the tail exponent of the distribution, or at least not in the regime of TTF and TTC values that can be explored within our dataset. In other words, even if the generalized Gamma distribution gives a (slightly) better fit in terms of likelihood, it is hard to estimate the tail exponent from the fitted parameters due to the slow convergence of the local exponent. On the contrary, the parameters estimated from the fit with functional form of Eq. (B1) give a better estimate of the tail exponent. In order to support this claim, we estimate independently the tail exponent by using the Hill estimator [52, 53]. In Fig. 9(right) we show the Hill plot of the time to fill of AZN with $\Delta = 0$. It is clear that the Hill estimator converges to a value which is much closer to 2.2 (as in the P_Z distribution) than to 4.5 (as in the P_G distribution).

In conclusion our analysis shows that, although the generalized Gamma distribution gives a slightly better *overall* fit of time to fill and time to cancel than our proposed form [Eq. (B1)], the parameters obtained from the fit of P_G suggest an unrealistic value of the tail exponent. On the contrary, our function P_Z allows us to both fit the data reasonably well *and* to obtain values of the tail exponent which are consistent with the Hill estimator.

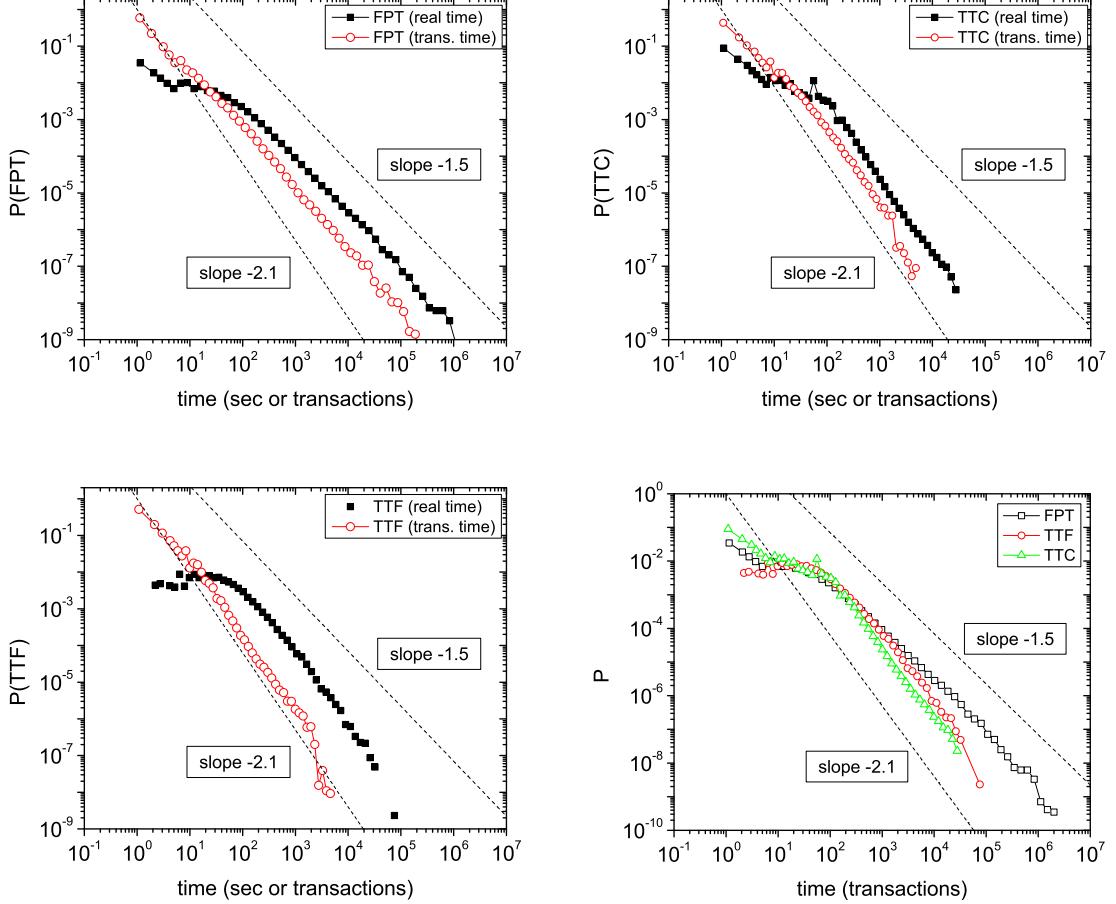


Figure 8: Comparison of the distribution of characteristic times in real time and transaction time, all plots are for GSK, $\Delta = 1$ tick. As a reference two power law decays with exponents 1.5 and 2.1 are also given. *Top left:* FPT, *Top right:* TTC, *Bottom left:* TTF, *Bottom right:* all three quantities, only in transaction time.

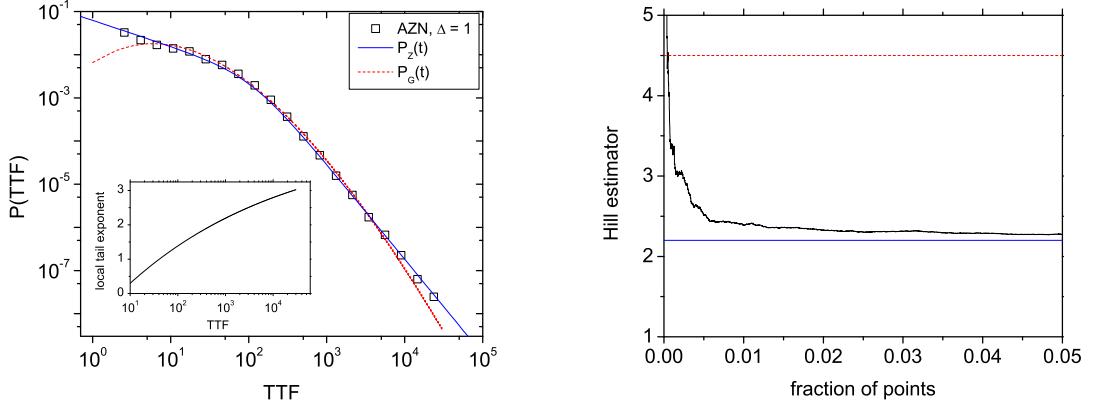


Figure 9: *Left:* Probability density function of the time to fill of limit orders with $\Delta = 0$ for AZN (boxes). We also show the Maximum Likelihood best fit according to the generalized Gamma distribution P_G of Eq. B2 (red dashed line) and to the functional form P_Z of Eq. B1 used in this paper (solid blue line). The inset shows the local tail exponent of the generalized Gamma distribution ($d[\log P_G(t)]/d[\log t]$) as a function of t . *Right:* Hill plot of the same data used in the left panel, showing the estimated tail exponent of the probability density function as a function of the fraction of points used in the estimation (black line). The dashed red line and the solid blue line are the values of the exponent obtained from the fit by P_G and P_Z , respectively.